PCA Method for Automated Detection of Mispronounced Words

Zhenhao Ge, Sudhendu R. Sharma, Mark J.T. Smith

School of Electrical and Computer Engineering
Purdue University, West Lafayette, IN, USA, 47907

arXiv:1602.08128v1 [cs.SD] 25 Feb 2016

ABSTRACT

This paper presents a method for detecting mispronunciations with the aim of improving Computer Assisted Language Learning (CALL) tools used by foreign language learners. The algorithm is based on Principle Component Analysis (PCA). It is hierarchical with each successive step refining the estimate to classify the test word as being either mispronounced or correct. Preprocessing before detection, like normalization and time-scale modification, is implemented to guarantee uniformity of the feature vectors input to the detection system. The performance using various features including spectrograms and Mel-Frequency Cepstral Coefficients (MFCCs) are compared and evaluated. Best results were obtained using MFCCs, achieving up to 99% accuracy in word verification and 93% in native/non-native classification. Compared with Hidden Markov Models (HMMs) which are used pervasively in recognition application, this particular approach is computational efficient and effective when training data is limited.

Keywords: Mispronunciation, PCA

1. INTRODUCTION

With the advent of technology, many Computer Assisted Language Learning (CALL) and Computer Assisted Pronunciation Training (CAPT) tools are available in the market to help people learn a new language. Learning a new language can be a very difficult. Without proper feedback, it can be very frustrating. Most CALL tools focus on teaching new words or sentences using a listen-and-repeat procedure. These tools use modern automatic speech recognition (ASR) algorithms and provide visual aids such as spectrograms and waveforms as feedback. However, these tools are not geared towards identifying specific mispronunciations and consequently are a poor substitute for a human instructor. The work discussed in this paper focuses on mispronunciation detection.

In the last two decades, a significant amount of research has been carried out in the field of mispronunciation detection. The majority of the mispronunciation detection studies[1–7] have extensively used ASRs based on statistical models, such as Hidden Markov Models (HMM). However, as García-Moral et al.[8] point out, for HMM based systems "large databases are required to warrant relevant and statistically reliable results". For mispronunciation detection, these systems have to train different models for correct pronunciations and mispronunciations. Recent works by Wang et al.,[3] and Harrison et al.[5] have shown improvement in detection accuracy, but these systems had to use extended recognition network by making use of cross-language phonological rules, i.e. rules that dictate how a learner's first language affects his/her pronunciation of a second language. The extended recognition networks require additional models to be trained in order to account for the cross-language rules. The use of large database and a large number of models makes HMM-based systems computationally complex and costly. Furthermore, performance often suffers when the training set is limited.

In this paper, we present a mispronunciation detection system based on Principal Component Analysis (PCA) and show it is computationally efficient compared to HMM's and performs well with limited training data. The approach used in this paper for detecting mispronunciation is to classify pronunciations, at both the word-level and syllable-level, as either acceptable or unacceptable.

Further author information: (Send correspondence to Zhenhao Ge)
Zhenhao Ge: E-mail: zge@purdue.edu, Telephone: 1 317 457 9348
Sudhendu R. Sharma: E-mail: sharmasr@purdue.edu, Telephone: 1 405 334 9992

2. PCA METHOD FOR PATTEN RECOGNITION

PCA is a well-known method that has been successfully employed in high-dimensional recognition problems. One of the popular applications of PCA is for face recognition which has been studied extensively over the past 20 years.[9–13] The fundamental idea behind PCA is encoding the most relevant information that distinguishes one pattern from another and eliminating dimensions with less information to reduce the computational cost.[14]

The mispronunciation detection work in this paper shares much in common with face recognition. Mispronunciation detection can be thought of as a pattern recognition problem, just like face recognition. Furthermore, Mel-Frequency Cepstral Coefficients (MFCCs) and spectrograms, which are the features that we are investigating, are also of high dimension. The similarities between mispronunciation detection and face recognition motivated our incorporating PCA-based face recognition techniques into our work. In the next section, we first review the application of PCA in face recognition, and then discuss the procedures of the PCA-based mispronunciation detection.

2.1 Review of Face Recognition using PCA

The PCA approach for face recognition requires all data (face images), training data $\mathcal{D}_{\text{train}}$, and testing data $\mathcal{D}_{\text{test}}$ from multiple classes to have the same size (say $N_1 \times N_2$). The method tends to work best when the faces in all the images are located at the same position within the image. These data matrices are then vectorized to N_1N_2-dimensional column vectors. When implementing a PCA-based detection/classification system, like face recognition, given M vectorized training faces $\Gamma_1, \Gamma_2, \ldots, \Gamma_M (M \ll N_1N_2)$ as the whole set of $\mathcal{D}_{\text{train}}$, let $\Psi = \frac{1}{M}\sum_{i=1}^{M}\Gamma_i$ be the mean face and let $\Phi_i = \Gamma_i - \Psi$ be the mean-shifted faces which form $A = [\Phi_1, \Phi_2, ..., \Phi_M]$. $C = AA^T$ is the covariance matrix of A, and can supply up to M eigenvectors to represent the eigenspace of A. Each face in A is a linear combination of these eigenvectors. If selecting the most significant M' eigenvectors to form a sub-eigenspace of training data $U = [u_1, u_2, \ldots, u_{M'}]$ $(N_1N_2 \times M')$, the representation of Φ_i in the M'-dimensional eigenspace is $\Omega_i = U^T\Phi_i$ and the projection of Φ_i in the eigenspace is $\hat{\Phi}_i = U\Omega_i$. The dimension of data is reduced from N_1N_2 to M' while preserving the most relevant information to distinguish faces.

After the eigenspace U for the training data A is computed, the following two steps are used for face recognition:[13]

1. Face detection: compute the Euclidean distance from test image Φ_j to its projection $\hat{\Phi}_j$ onto the eigenspace U,
$$e_{\text{dfes}} = \|\Phi_j - \hat{\Phi}_j\| \ , \tag{1}$$
where e_{dfes} is called "distance from eigenspace". If $e_{\text{dfes}} < T_d$, where T_d is a detection threshold, the test image is verified to be a face and step 2 below is used for face classification.

2. Face classification: compute the minimum Euclidean distance between Ω_j and Ω_k in Equation (2), where Ω_j is the eigenspace representation of the test face Φ_j and Ω_k is the averaged eigenspace representation of faces in class k, $k = 1, 2, \ldots, K$, where K is the number of classes.
$$e_{\text{dies}} = \min \|\Omega_j - \Omega_k\| \qquad k = 1, 2, \ldots, K \tag{2}$$
where e_{dies} is called "distance within eigenspace". If $e_{\text{dies}} < T_c$, where T_c is the classification threshold, then the test image is classified as a face belonging to that class; otherwise, a new class of faces is claimed. Figure 1 illustrates the difference between e_{dfes} and e_{dies}.

2.2 Procedure for Mispronunciation Detection Using PCA

In this work, the training and testing data consist of two disjointed classes of speech samples. The first class consists of samples from native speakers, all of which have correct pronunciations. The second class contains samples from non-native speakers with possible mispronunciations. Given an input test sample, the procedure used for mispronunciation detection can be divided into 3 steps:

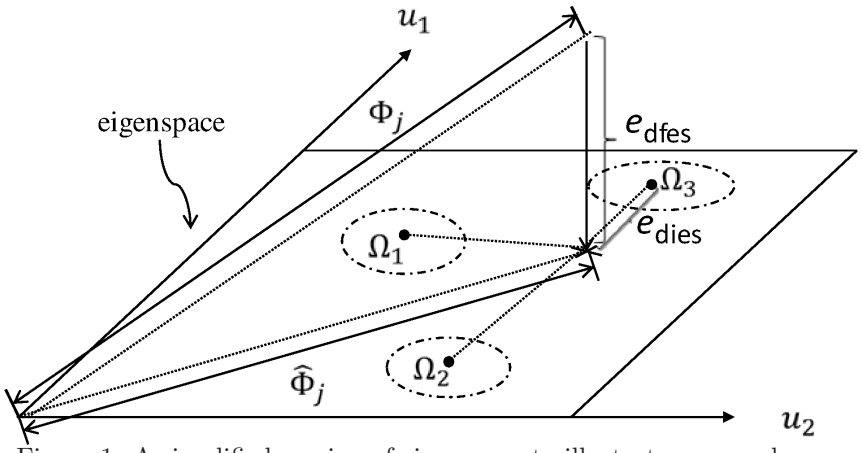

Figure 1: A simplified version of eigenspace to illustrate e_{dfes} and e_{dies}

1. Word verification: compute e_{dfes} from the test sample Φ_j to the eigenspace U_{All} constructed from all samples in $\mathcal{D}_{\text{All.train}}$, where $\mathcal{D}_{\text{All.train}}$ includes all native $\mathcal{D}_{\text{N.train}}$ and non-native $\mathcal{D}_{\text{NN.train}}$ samples. If $e_{\text{dfes}} < T_d$, where T_d is the threshold for word verification, then the test sample is considered "verified" and we proceed to step 2. Otherwise the test sample is "rejected" and the detection process stops.

2. Native/non-native (N/NN) classification: compute e_{dfes} from the test sample Φ_j to the eigenspace U_{N} constructed only from $\mathcal{D}_{\text{N.train}}$. If $e_{\text{dfes}} < T_c$, then classify the test sample as native and stop. Otherwise, classify it as non-native and proceed to step 3.

3. Syllable-level mispronunciation detection: Divide the non-native test samples into syllables $\Phi_{jk}, k = 1, 2, \ldots, K$, where K is the number of syllables for that test sample. Similarly divide native samples $\mathcal{D}_{\text{N.train}}$ into syllables and call the k^{th} syllable $D_{\text{N}k\text{.train}}$. For each test syllable Φ_{jk}, compute its e_{dfes} to the eigenspace $U_{\text{N}k}$ constructed from k^{th} syllable of native samples $\mathcal{D}_{\text{N}k\text{.train}}$. If $e_{\text{dfes}} > T_k$, then classify the test syllable as mispronounced, otherwise classify it as correctly pronounced.

Word verification is very similar to face detection, since both cases use the distance metric e_{dfes} of the eigenspace U_{All} which is constructed from all classes of samples. However, N/NN classification is different from face classification. The main difference is that in face classification e_{dies} is used as the metric to determine the class of the "test" face, where as in mispronunciation detection we use e_{dfes} as the metric to do the N/NN classification. In face classfication, the various classes of faces are well defined. If e_{dies} of a "test face" lies within the predefined threshold of some class k, then it is classified as face k. Otherwise, it is claimed to be a "new face" that does not match any of the faces. The "new face" scenario for face classification is shown in Figure 2a, where Ω_j is the "test" face. However, in N/NN classification, there are only two classes and any region outside the native class belongs to the non-native class. So the misclassification shown in Figure 2b will occur if e_{dies} is used to determine the class of Ω_j. In Figure 2b, Ω_j is in the non-native class but it is closer to native class. Thus, instead of using distance e_{dies} within eigenspace U_{All}, N/NN classification uses e_{dfes} to measure the distance from the eigenspace U_{N}, which is a sub-eigenspace of U_{All}.

The syllable-level mispronunciation detection is an extension of the N/NN classification applied to syllables. Each "test syllable" is treated like a "test word" in step 2, and mispronunciation can be detected in the syllables using the same N/NN classification approach.

3. SYSTEM DESIGN AND IMPLEMENTATION

Having reviewed the application of PCA in face recognition and how it can be applied to mispronunciation detection, the following section will discuss the implementation of PCA-based mispronunciation detection in detail, including database construction, data pre-processing, feature selection, eigenspace training, and detection threshold optimization.

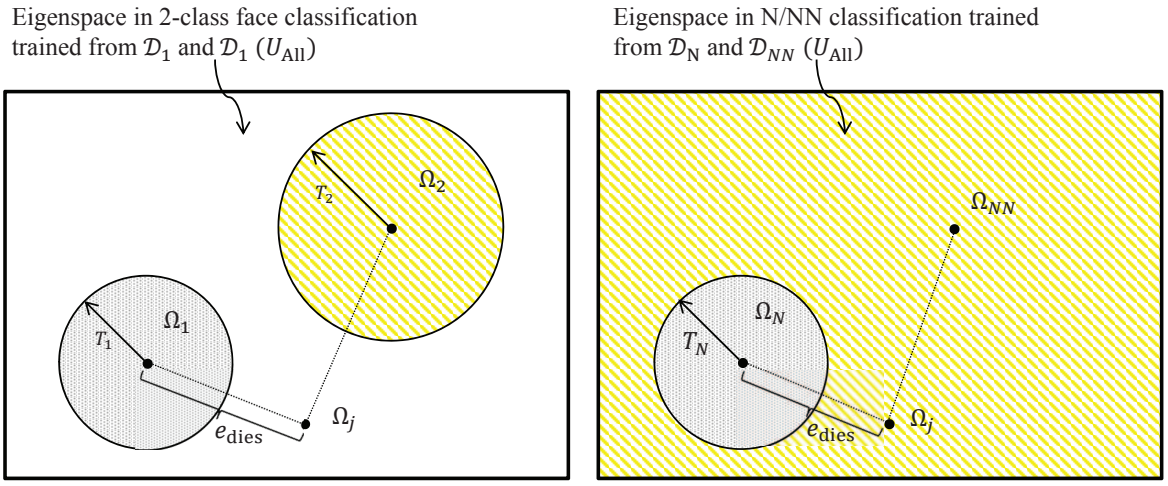

(a) case 1 (b) case 2

Figure 2: Possible distributions of native and non-native samples in the eigenspace U_{All}

3.1 Database Construction

Although the algorithm is not language specific, Spanish was chosen for this work. The database used in this work is relatively small and contains only 10 Spanish words listed in Table 1. These words were selected by language experts to cover a variety of common mispronunciations exhibited by American speakers learning Spanish. There were 13 male speakers, 7 of them native speakers and 6 non-native. Each speaker repeated each of the 10 words 5 times.

Because of the limited size of the database, the Leave-One-Out method was used for training and testing and is further discussed in Section 3.4.

Table 1: List of Spanish words and syllables

No.	Words	Syllables	No.	Words	Syllables
1	jamaica	/ja/, /mai/, /ca/	6	tordurados	/tor/, /tu/, /ra/, /do/, /s/
2	tres	/tr/, /e/, /s/	7	accidente	/ac/, /ci/, /den/, /te/
3	gemelas	/ge/, /me/, /la/, /s/	8	construccion	/cons/, /truc/, /cion/
4	hierro	/hie/, /rro/	9	puertorriquena	/puer/, /tor/, /ri/, /que/, /na/
5	pala	/pa/, /la/	10	aire	/ai/, /re/

3.2 Data Pre-processing and Feature Extraction

The PCA method requires centralized and uniform-size input features for training and testing. For centralization, periods of silence before and after the actual speech segment were removed using voiced/unvoiced detection. To make all samples uniform in size, the samples were time-scale modified to match the average duration of the training data. Furthermore, to improve the detection performance, the background noise was suppressed and the amplitude of each sample was normalized to unity.

Spectrograms and MFCCs were chosen as input features to the detection system. These features were computed for frames with window size 25ms (using a Hamming window) and 15 ms overlap between frames. The spectrogram feature space is 50-dimensional with each dimension spanning 320 Hz to 16 KHz. The MFCC feature space was 13-dimensional representing the first 13 Mel-scale cepstral coefficients.

Table 2: Data and eigenspace comparison of 3 mispronunciation detection steps

	Step 1	**Step 2**	**Step 3**
Eigenspace	$U_{\text{All}}(W_{i^*})$	$U_{\text{N}}(W_{i^*})$	$U_{\text{N}}(W_{i^*}, k)$
Data used to train eigenspace	$\mathcal{D}_{\text{All}}(W_{i^*})$	$\mathcal{D}_{\text{N}}(W_{i^*})$	$\mathcal{D}_{\text{N}}(W_{i^*}, k)$
Class 1 data used to find e^1_{dfes}	$\mathcal{D}_{\text{All}}(W_{i^*})$	$\mathcal{D}_{\text{N}}(W_{i^*})$	$\mathcal{D}_{\text{N}}(W_{i^*}, k)$
Class 2 data used to find e^2_{dfes}	$\mathcal{D}_{\text{All}}(W_i)$,$i = 1, 2, ..., L$, but $i \neq i^*$	$\mathcal{D}_{\text{NN}}(W_{i^*})$	$\mathcal{D}_{\text{NN}}(W_{i^*}, k)$

Note: U denotes eigenspace, $\mathcal{D}$ denotes data, All denotes both native and non-native included, N denotes native, NN denotes non-native, L denotes the number of different words, (W_{i^*}, k) denotes k^{th} syllable in word i^*, and $k = 1, 2, ..., K(i^*)$ where $K(i^*)$ is the number of syllables in word $W(i^*)$

3.3 Eigenspace Training and Detection Threshold Optimization

As discussed in Section 2.2, these three steps are: (a) word verification; (b) N/NN classification; and (c) syllable-level mispronunciation detection. The detection system runs on a word-by-word basis. For each word, W_{i^*}, $i^* = 1, 2, ..., 10$, each step (a), (b), and (c) undergoes the following 3 phases of training.

- Phase 1: Train to determine eigenspace U.
- Phase 2: Compute two sets of distances e^1_{dfes} and e^2_{dfes}, where e^1_{dfes} corresponds to the distances from "class 1" samples to the eigenspace U and e^2_{dfes} corresponds to the distance from "class 2" samples to U.
- Phase 3: Find an optimal detection threshold T that separates these two sets of distances.

Even though the 3 training phases are the same for the 3 steps (a), (b), and (c), each step has a different trained eigenspace, uses different class 1 and class 2 data, and generates different thresholds. Table 2 summarizes the difference in the 3 training phases of the 3 steps. In Table 2, the data used in training to obtain the eigenspace (2^{nd} row) and to compute e^1_{dfes} (3^{rd} row) are slightly different. Even though they share the same notation in the table, they represent two disjoint subsets within the class 1 data. Furthermore, different eigenspaces are trained using the Leave-One-Out approach. In this approach, samples corresponding to one speaker are left out. The samples from remaining speakers are used to train the eigenspace, and the samples from the left out speaker are used to find e^1_{dfes}. This is repeated for all speakers belonging to class 1. As mentioned, this approach leads to several trained eigenspaces. The e^2_{dies} distances are then computed for class 2 data to those different trained eigenspaces and averaged for each speaker in class 2. The distributions formed by e^1_{dfes} and e^2_{dfes}, an example shown in Figure 3, are then used to find the optimal threshold. By convention,[15] the dimension of the eigenspace in each step is determined by selecting eigenvectors that represent 80% variance of the total principle components. For example, the dimension of the eigenspace used to determine e_{dfes} in Figure 3 is 18.

The main goal of the training process is to find the "optimal threshold" to detect mispronunciations in the test samples. In each step, since the test samples come from two classes of data which are supposed to be separated, the optimal threshold T should be the one that separates these two classes best. This is done using Bayes rule by finding the optimal threshold to separate two sets of distances e^1_{dfes} for class 1 and e^2_{dfes} for class 2 data. Let random variables X_1 and X_2 represent e^1_{dfes} and e^2_{dfes} and assume they are both Gaussian distributed with prior probabilities $P(\omega_1)$ and $P(\omega_2)$, where $\omega_{1,2}$ denote the classes of these two groups. The classification error can be computed by Equation (3) using Bayes rule

$$P(\text{error}) = \int_{x^*}^{\infty} p(x \mid \omega_1)P(\omega_1)\mathrm{d}x + \int_{-\infty}^{x^*} p(x \mid \omega_2)P(\omega_2)\mathrm{d}x \tag{3}$$

where the optimal threshold x^* or T can be found by computing the discriminant function $g(x)$ at $g(x) = 0$ where

$$g(x) = p(x \mid \omega_1)P(\omega_1) - p(x \mid \omega_2)P(\omega_2) \tag{4}$$

Figure 3 illustrates the numerical distribution of e^1_{dfes} (top line) and e^2_{dfes} (bottom line) as an example of word verification step for the word `tres`. The distances are averaged at 5 repetitions for each speaker. Figure 4 plots the corresponding theoretical Gaussian distribution in order to obtain a more reliable optimal threshold. MFCCs were used as the feature set here. Note that since there are 9 words from class 2 and only 1 word from class 1, distance e^1_{dfes} has been repeated 9 times to make them comparable to the e^2_{dfes}. The optimal threshold T_d can be found by computing the minimum error rate of the theoretical distribution (Gaussian distribution assumed) shown in Figure 4. By Bayes rule, the threshold is optimal when it provides the theoretical minimum classification error P(error) (0.030% in this case). The P(error) computed using spectrograms is similar but a little higher (0.447%).

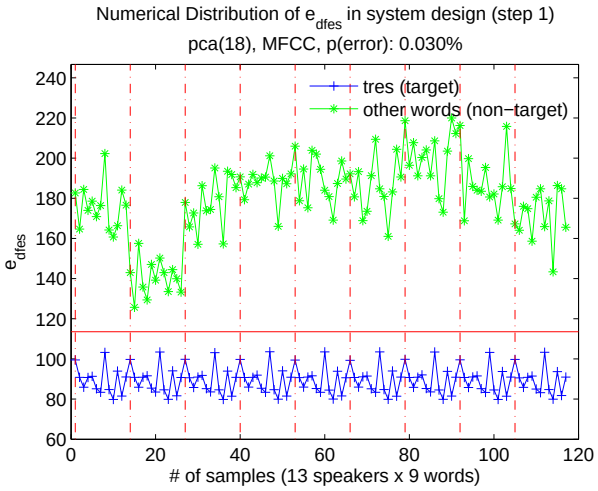
Numerical Distribution of e_dfes in system design (step 1)
pca(18), MFCC, p(error): 0.030%
tres (target)
other words (non−target)
e_dfes
# of samples (13 speakers x 9 words)

Figure 3: Numerical distribution of class 1 and class 2 data in word verification (target word: `tres`, feature: MFCCs)

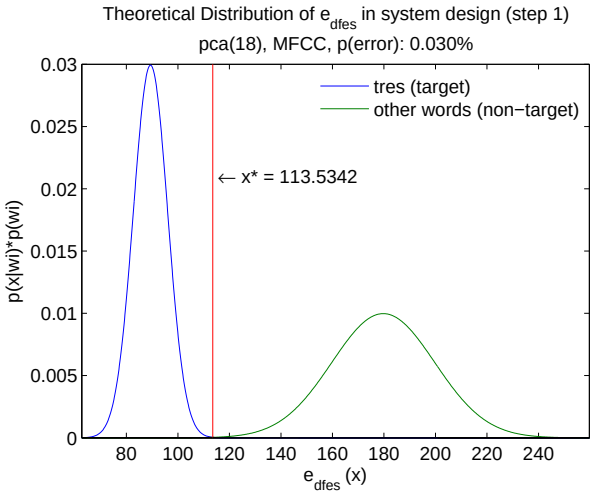

Figure 4: Theoretical distribution of class 1 and class 2 data in word verification (target word: `tres`, feature: MFCCs)

Figure 5 illustrates the native/non-native distribution of data in N/NN classification step for word `pala`. The P(error) obtained using MFCCs (Figure 6) and spectrograms are 0.001% and 11.65% respectively. Experimental results shows that MFCCs are much better than spectrograms in separating data.

4. SYSTEM TESTING AND RESULTS

After eigenspace training and detection threshold optimization, the mispronunciation detection system is built up. The following section presents the system testing results in each step.

4.1 Leave-One-Out Training and Testing

Because of the relatively small size of the database, the Leave-One-Out (LOO) method is used for training and testing. Traditionally in LOO, all but one sample is used in training and the left out sample is used for testing. In our case, samples belonging to one speaker (i.e. all 5 repetitions) are left out for testing and all samples belonging to all other speakers are used in system training, which includes the 3 phases discussed in the Section 3.3: eigenspace construction (U), 2-class distance measurement (e^1_{dfes},e^2_{dfes}), and detection threshold optimization (T_d, T_c and T_k).

For example, in the word verification step, the goal is to verify whether or not the "test word" is the target word W_{i^*}. The number of samples available for training and testing in this step is 650 (10 different words × 13 speakers × 5 repetitions). Using the LOO method, 5 repetitions of word W_{i^*} from one speaker are left out for

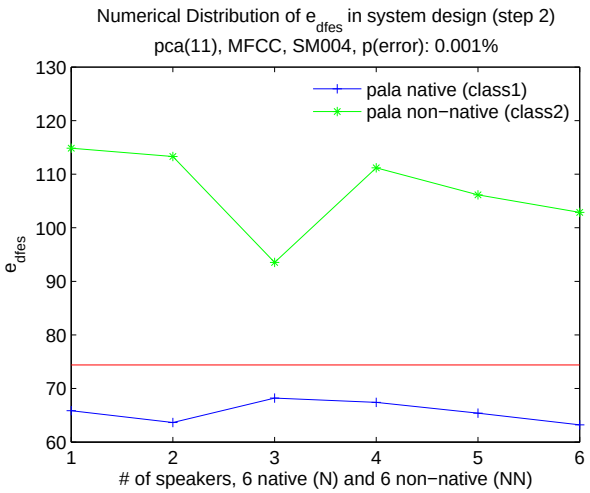
Numerical Distribution of e dfes in system design (step 2)
pca(11), MFCC, SM004, p(error): 0.001%
pala native (class1)
pala non−native (class2)
e dfes
# of speakers, 6 native (N) and 6 non−native (NN)

Figure 5: Numerical distribution of e_{dfes} of the class 1 and class 2 data in N/NN classification (target word: `pala`, feature: MFCCs)

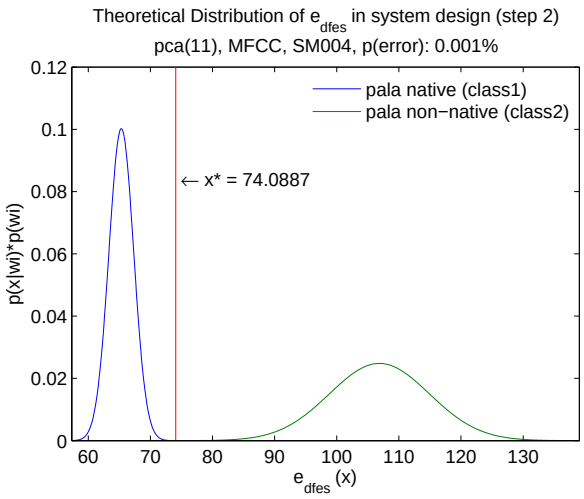
Theoretical Distribution of e dfes in system design (step 2)
pca(11), MFCC, SM004, p(error): 0.001%
pala native (class1)
pala non−native (class2)
← x* = 74.0887
p(x|wi)*p(wi)
e dfes (x)

Figure 6: Theoretical distribution of e_{dfes} of the class 1 and class 2 data in N/NN classification (target word: `pala`, feature: MFCCs)

testing and the remaining 645 samples are used to train the system and obtain the optimal threshold T_d. This is repeated for each speaker. At the end, there are 130 trained system/threshold combinations and 5 test samples for each system to be validated.

In the native/non-native classification, the number of samples available for training and testing is 65 (13 speakers $\times$ 5 repetitions of the target word). Five repetitions from one speaker are left out for testing and the remaining 60 samples are used to train the system and obtain the optimal threshold T_c. At the end, there are 13 system/threshold combinations and 5 test samples for each system to be validated.

4.2 Results of Word Verification and N/NN Classification

Compared with the theoretical error rate in Equation (3), the performance of the mispronunciation detection system is measured by the numerical error rate P_e

$$P_e = \frac{N_{e1} + N_{e2}}{N_1 + N_2} \ , \tag{5}$$

where N_1, N_2 are the number of test samples from class 1 and 2, and N_{e1}, N_{e2} are the number of misclassified samples from each class.

In word verification, the error rate P_e is always below 3% and 7% for MFCCs and spectrograms respectively. In N/NN classification, the performance based on HMMs using MFCCs is also compared along with the PCA method. Figure 7 and 8 show the results of the Leave-One-Out method applied to the word `aire` using PCA and HMMs. The 13 columns in the figure represent the 13 speakers of the word `aire`. For each column (speaker), there are 5 samples, which are compared against the threshold. The bottom line corresponds to the 7 native speakers and the top line corresponds to the 6 non-native speakers. The thresholds during each LOO iteration vary slightly because of small differences in the training database during each trial. These variations diminish as the database size increases. For the PCA method using MFCCs, there is 1 sample from both native speaker 4 and 6 in Figure 7 that are slightly above the threshold and misclassified as non-native and all the rest from both classes are correctly classified. A similar situation is illustrated in Figure 8 with slightly higher P_e. Complete word verification and N/NN classification results of all 10 words are provided in Table 3.

In Table 3, MFCCs are shown to perform much better than spectrograms using PCA, especially in step 2, N/NN classification. The PCA method performs slightly better on average than HMMs in this data-limited case. However, general parameters of HMMs are used and they are not optimized for this specific system. With respect to the computational cost, training an eigenspace is almost 10^3 times faster than training HMMs (5 ms

Table 3: Error rate P_e in word verification and N/NN classification

Steps	Methods -Features	Words W_1	W_2	W_3	W_4	W_5	W_6	W_7	W_8	W_9	W_{10}	Max	Min	Avg.
1	PCA-MFCCs	2.04%	0.03%	0.63%	0.07%	0.19%	0.18%	0.37%	1.04%	0.08%	1.43%	2.04%	0.03%	**0.61%**
	PCA-Spec.	6.70%	0.45%	1.56%	2.67%	5.40%	3.47%	3.29%	5.24%	2.22%	6.67%	6.70%	0.45%	**3.77%**
2	PCA-MFCCs	12.31%	10.77%	1.54%	12.31%	7.69%	7.69%	6.15%	9.23%	4.62%	3.08%	12.31%	1.54%	**7.54%**
	PCA-Spec.	24.62%	16.92%	27.69%	21.54%	18.46%	18.46%	10.77%	18.85%	13.85%	15.38%	27.69%	10.77%	**18.15%**
	HMM-MFCCs	3.08%	10.77%	6.51%	9.23%	10.77%	12.31%	9.23%	7.69%	4.62%	6.154%	12.31%	3.08%	**8.04%**

vs. 5 s), while in sample testing, the PCA method is also 60-80 times faster (about 4 ms vs. 250 ms) than HMMs depending on the length of the word.

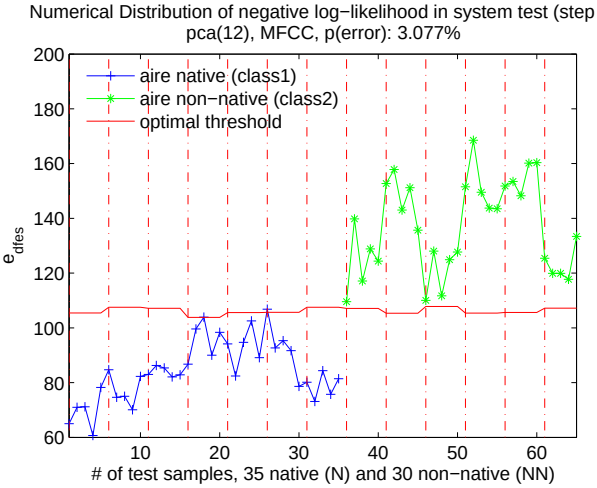
Numerical Distribution of negative log−likelihood in system test (step 2)
pca(12), MFCC, p(error): 3.077%
aire native (class1)
aire non−native (class2)
optimal threshold
e_dfes
# of test samples, 35 native (N) and 30 non−native (NN)

Figure 7: Numerical distribution of e_{dfes} of the test samples in N/NN classification using PCA (target word: `aire`, feature: MFCCs)

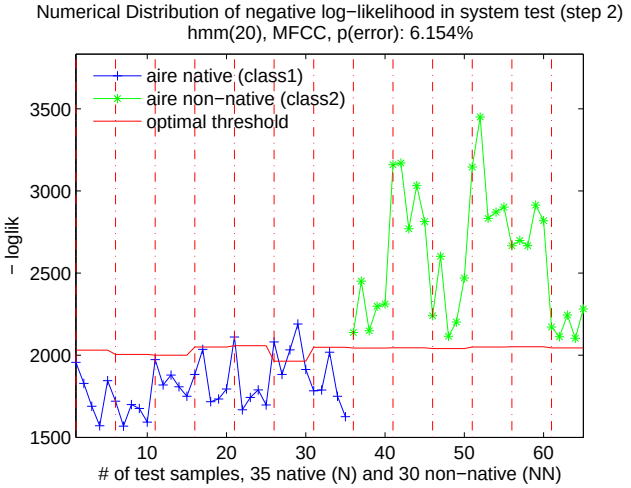
Numerical Distribution of negative log−likelihood in system test (step 2)
hmm(20), MFCC, p(error): 6.154%
aire native (class1)
aire non−native (class2)
optimal threshold
− loglik
# of test samples, 35 native (N) and 30 non−native (NN)

Figure 8: Numerical distribution of negative log-likelihood of the test samples in N/NN classification using HMM (target word: `aire`, feature: MFCCs)

4.3 Results of Syllable-Level Mispronunciation Detection

For syllable-level mispronunciation detection, an approach similar to the N/NN classification is followed, except that it is done at a syllable level. Though in real applications, only samples that are classified as non-native in the N/NN classification would be processed by the 3$^{\text{rd}}$ step, to make testing results comparable and unbiased, all test samples including native samples are also used here for evaluation.

There is a major difference between this step and the previous two steps. This is because the assumption that two classes of data in eigenspace training and e_{dfes} computing are separated is no longer valid, namely, some syllables may be pronounced well enough that can not be used to distinguish native and non-native. If this is the case, the threshold obtained using Bayes rules is biased and moves towards the native class, which dramatically increases the classification error rate P_e. Thus, P_e cannot be used to measure the performance of mispronunciation detection. Instead, it serves as the similarity measurement of the syllables pronounced by both native and non-native classes. By dividing the total error rate P_e into False Negative Rate (FNR) and False Positive Rate (FPR) in Equation (6), it is easy to find that with a relatively low FNR, higher FPR indicates better pronunciation of the syllable, while lower FPR shows the problematic syllable that one should pay attention to.

$$\text{FNR} = \frac{N_{e1}}{N_1} \quad \text{and} \quad \text{FPR} = \frac{N_{e2}}{N_2} \tag{6}$$

Figure 9 and 10 illustrate two examples where "good" (FPR = 80%) and "bad" (FPR = 23.3%) mispronunciations are detected. Table 4 shows FNR/FPR for each syllable and highlights all FPR $\leq$ 30%. From Table 4, common mispronunciation problems[16] like vowel reduction, aspiration, linkage and stress are successfully detected in the test database. Some of them are listed below:

- Vowel Reduction: /pa/ in `pala`; /den/ in `accidente`; /ie/ in `arie`
- Aspiration: /pa/ in `pala`; /puer/ in `puertorriquena`
- Linkage: /tr/ in `tres`; /cons/ in `construccion`; /na/ in `puertorriquena`
- Stress: /ci/ in `accidente`; /cons/ and /cion/ in `construccion`

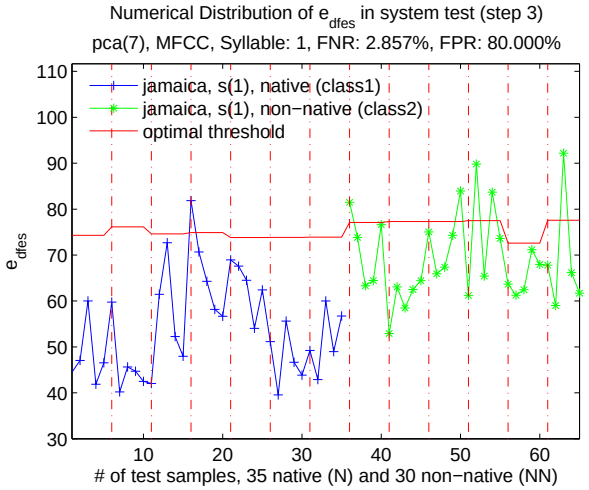
Numerical Distribution of e_dfes in system test (step 3)
pca(7), MFCC, Syllable: 1, FNR: 2.857%, FPR: 80.000%
jamaica, s(1), native (class1)
jamaica, s(1), non−native (class2)
optimal threshold
e_dfes
# of test samples, 35 native (N) and 30 non−native (NN)

Figure 9: Numerical distribution of e_{dfes} of the test samples in mispronunciation classification using PCA (target syllable: `jamaica`, /ja/, feature: MFCCs)

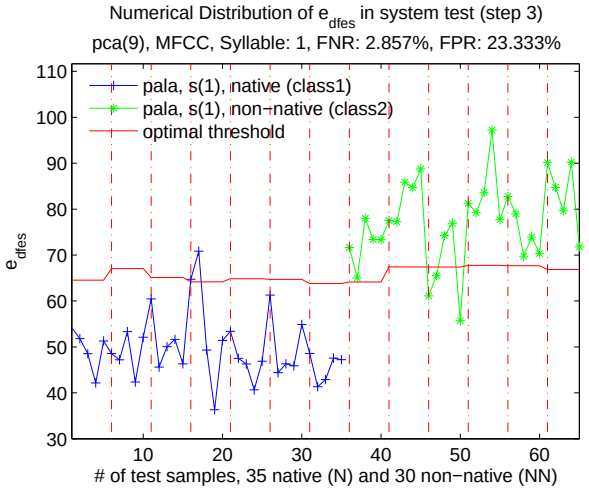
Numerical Distribution of e_dfes in system test (step 3)
pca(9), MFCC, Syllable: 1, FNR: 2.857%, FPR: 23.333%
pala, s(1), native (class1)
pala, s(1), non−native (class2)
optimal threshold
e_dfes
# of test samples, 35 native (N) and 30 non−native (NN)

Figure 10: Numerical distribution of e_{dfes} of the test samples in mispronunciation classification using PCA (target syllable: `pala`, /pa/, feature: MFCCs)

5. CONCLUSION AND FUTURE WORK

The testing results in Section 4 show that PCA can be a computationally efficient approach to detect mispronunciations, even when the training and testing database is limited. MFCCs were shown to outperform spectrograms. Compared with HMMs, the PCA method is much faster and achieves comparable results for the database used in this paper.

For future work, the threshold in each step should be optimized on a larger database to improve robustness. This is especially important for optimizing the threshold in the syllable-level detection, since subtle differences that reside in syllables require more data to differentiate.

Because of the constraints of the PCA method in data separation (discriminant information may stay in the less significant components), further investigation on the distributions of multiple types of mispronunciations is warranted. Furthermore, hybrid methods based on PCA, LDA (Linear Discriminant Analysis) and ICA (Independent Component Analysis) should also be considered.

Table 4: FNR and FPR in mispronunciation detection

Words	Syllables	FNR*	FPR**	Words	Syllables	FNR	FPR
jamaica	/ja/	2.86%	80.00%	pala	/pa/	2.86%	**23.33%**
	/mai/	0.00%	80.00%		/la/	0.00%	80.00%
	/ca/	11.43%	53.33%	accidente	/ac/	5.71%	53.33%
tres	/tr/	5.71%	**16.67%**		/ci/	2.86%	**26.67%**
	/e/	0.00%	73.33%		/den/	5.71%	**10.00%**
	/s/	11.43%	36.67%		/te/	8.57%	66.67%
gemelas	/ge/	5.71%	40.00%	construccion	/cons/	0.00%	**30.00%**
	/me/	2.86%	43.33%		/truc/	2.86%	40.00%
	/la/	2.86%	80.00%		/cion/	8.57%	**30.00%**
	/s/	5.71%	60.00%	puertorriquena	/puer/	5.71%	**30.00%**
hierro	/hie/	2.86%	56.67%		/tor/	2.86%	80.00%
	/rro/	8.57%	40.00%		/ri/	0.00%	70.00%
torturados	/tor/	8.57%	**30.00%**		/que/	2.86%	70.00%
	/tu/	2.86%	43.33%		/na/	2.86%	**26.67%**
	/ra/	8.57%	70.00%	aire	/ai/	2.86%	33.33%
	/do/	17.14%	73.33%		/re/	0.00%	**30.00%**
	/s/	8.57%	86.67%				

*FNR: False Negative Rate; **FPR: False Positive Rate

REFERENCES

1. Franco, H., Abrash, V., Precoda, K., Bratt, H., Rao, R., Butzberger, J., Rossier, R., and Cesari, F., "The sri eduspeak system: Recognition and pronunciation scoring for language learning," in [*Proc. of InSTIL*], 123–128 (2000).
2. Ronen, O., Neumeyer, L., and Franco, H., "Automatic detection of mispronunciation for language instruction," in [*Proc. of Eurospeech*], (1997).
3. Wang, L., Feng, X., and Meng, H. M., "Mispronunciation detection based on cross-language phonological comparisons," in [*IEEE IET International Conference on Audio, Language and Image Processing*], 307–311 (2008).
4. Harrison, A., Lau, W., Meng, H., and Wang, L., "Improving mispronunciation detection and diagnosis of learners' speech with context-sensitive phonological rules based on language transfer," in [*Interspeech*], 2787–2790 (2008).
5. Harrison, A., Lo, W., Qian, X., and Ming, H., "Implementation of an extended recognition network for mispronunciation detection and diagnosis in computer-assisted pronunciation training," in [*SLaTE*], 45–48 (2009).
6. Wei, S., Hu, G., Hu, Y., and Wang, R.-H., "A new method for mispronunciation detection using support vector machine based on pronunciation space models," *Speech Communication* **51**(10), 896 – 905 (2009).
7. Qian, X., Meng, H., and Soong, F., "Capturing l2 segmental mispronunciations with joint-sequence models in computer-aided pronunciation training (CAPT)," in [*International Symposium on Chinese Spoken Language Processing (ISCSLP)*], (2010).
8. García-Moral, A. I., Solera-Ureña, R., Peláez-Moreno, C., and de María, F. D., "Data balancing for efficient training of hybrid ANN/HMM automatic speech recognition systems," *IEEE Transactions on Audio, Speech & Language Processing* **19**(3), 468–481 (2011).
9. Sirovich, L. and Kirby, M., "A low-dimensional procedure for the characterization of human faces," *The Journal of the Optical Society of America* **4**, 519–524 (1987).

10. Kirby, M. and Sirovich, L., "Application of the karhunen-loeve procedure for the characterization of human faces," *IEEE Trans. on Pattern Analysis and Machine Intelligence* **12**(1), 103–107 (1990).
11. Belhumeur, P., Hespanha, J., and Kriegman, D., "Eigenfaces vs. fisherfaces: Recognition using class specific linear projection," *IEEE Trans. on Pattern Analysis and Machine Intelligence* **19**(7), 771–720 (1997).
12. Mandal, T. and Q.M.J.Wu, "Face recognition using curvelet based pca," in [*Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*], 1–4 (2008).
13. Turk, M. and Pentland, A., "Eigenfaces for recognition," *Journal of Cognitive Neuroscience* **3**(1), 71–86 (1991).
14. Shlens, J., "A tutorial on principal component analysis," tech. rep., Institute for Nonlinear Science, UCSD (2005).
15. Batch, S. A., "Principal variance component analysis," (August 2010). `http://www.niehs.nih.gov/research/resources/software/pvca`.
16. Hammond, R. M., [*The Sounds of Spanish: Analysis and Application with special reference to American English*], Cascadilla (2001).